Positivity of the English language

Isabel M. Kloumann,[1,*] Christopher M. Danforth,[2,†] Kameron Decker Harris,[2,‡] Catherine A. Bliss,[2,§] and Peter Sheridan Dodds[2,¶]

[1] *Center for Applied Mathematics, Cornell University, Ithaca, NY 14853*
[2] *Department of Mathematics and Statistics, Center for Complex Systems, & the Vermont Advanced Computing Center, University of Vermont, Burlington, VT, 05401*

(Dated: August 29, 2011)

Within the last million years, human language has emerged and evolved as a fundamental instrument of social communication and semiotic representation. People use language in part to convey emotional information, leading to the central and contingent questions: (1) What is the emotional spectrum of natural language? and (2) Are natural languages neutrally, positively, or negatively biased? Previous findings are mixed: suggestive evidence of a positive bias has been found in small samples of English words [1–3], framed as the Pollyanna Hypothesis [3] and Linguistic Positivity Bias [1], while the experimental elicitation of emotional words has instead found a strong negative bias [4]. Here, we report that the human-perceived positivity of over 10,000 of the most frequently used English words exhibits a clear positive bias. More deeply, we characterize and quantify distributions of word positivity for four large and distinct corpora, demonstrating that their form is surprisingly invariant with respect to frequency of word use.

While we regard ourselves as social animals, we have a history of actions running from selfless benevolence to extreme violence at all scales of society, and we remain scientifically and philosophically unsure as to what degree any individual or group is or should be cooperative and pro-social. Traditional economic theory of human behavior, for example, assumes that people are inherently and rationally selfish—a core attribute of *homo economicus*—with the emergence of global cooperation thus rendered a profound mystery [5, 6]. Yet everyday experience and many findings of psychology, behavioral economics, and neuroscience indicate people favour seemingly irrational heuristics [7, 8] over strict rationality as exemplified in loss-aversion [9], confirmation bias [10], and altruistic punishment [11]. Religions and philosophies similarly run the gamut in prescribing the right way for individuals to behave, from the universal non-harming advocated by Jainism, Gandhi's call for non-violent collective resistance, and exhortations toward altruistic behavior in all major religions, to arguments for the necessity of a Monarch [12], the strongest forms of libertarianism, and the "rational self-interest" of Ayn Rand's Objectivism [13].

In taking the view that humans are in part storytellers—*homo narrativus*—we can look to language itself for quantifiable evidence of our social nature. How is the structure of the emotional content rendered in our stories, fact or fiction, and social interactions reflected in the collective, evolutionary construction of human language?

To test the overall positivity of the English language, and in contrast to previous work [2, 4, 14], we chose words based solely on frequency of use, the simplest and most impartial gauge of word importance. We focused on measuring happiness, or psychological valence [15], as it represents the dominant emotional response [16, 17]. With this approach, we examined four large-scale text corpora (see Tab. I for details):

1. Twitter,
2. The Google Books Project (English),
3. The New York Times, and
4. Music lyrics.

These corpora, which we will refer to as TW, GB, NYT, and ML, cover a wide range of written expression including broadcast media, opinion, literature, songs, and public social interactions ([18]).

We took the top 5000 most frequently used words from each corpus, and merged them to form a resultant list of 10,222 unique words. We then used Mechanical Turk [19, 25] to obtain 50 independent evaluations per word on a 1 to 9 integer scale, asking participants to rate their happiness in response to each word in isolation (1 = least happy, 5 = neutral, and 9 = most happy [14, 24]).

We computed the average happiness score for each word, obtaining sensible results that showed excellent agreement with previous studies for smaller word sets [14, 19]. The highest and lowest scores were h_{avg}('laughter')=8.50 and h_{avg}('terrorist')=1.30, with expectedly neutral words averaging near 5, e.g., h_{avg}('the')=4.98 and h_{avg}('it')=5.02.

In Fig. 1, we show distributions of word happiness for the four corpora (black curves). We see each distribution is unimodal and positively skewed, with a clear abundance of positive words ($h_{\text{avg}} > 5$, yellow shade) over negative ones ($h_{\text{avg}} < 5$, gray shade). In order, the percentages of positive words are 72.00% (TW), 78.80%

*Electronic address: isabel.kloumann@uvm.edu
†Electronic address: chris.danforth@uvm.edu
‡Electronic address: kameron.harris@uvm.edu
§Electronic address: catherine.bliss@uvm.edu
¶Electronic address: peter.dodds@uvm.edu

Corpus (Abbreviation):	Date range	# Words	# Texts	Reference
Twitter (TW)	9/9/2008 to 3/3/2010	9.07×10^{9}	8.21×10^{8} tweets	[19, 20]
Google Books Project, English (GB)	1520 to 2008	3.61×10^{11}	3.29×10^{6} books	[21, 22]
The New York Times (NYT)	1/1/1987 to 6/30/2007	1.02×10^{9}	1.8×10^{6} articles	[23]
Music lyrics (ML)	1960 to 2007	5.86×10^{7}	2.95×10^{5} songs	[24]

TABLE I: Details of the four corpora we examined for positivity bias.

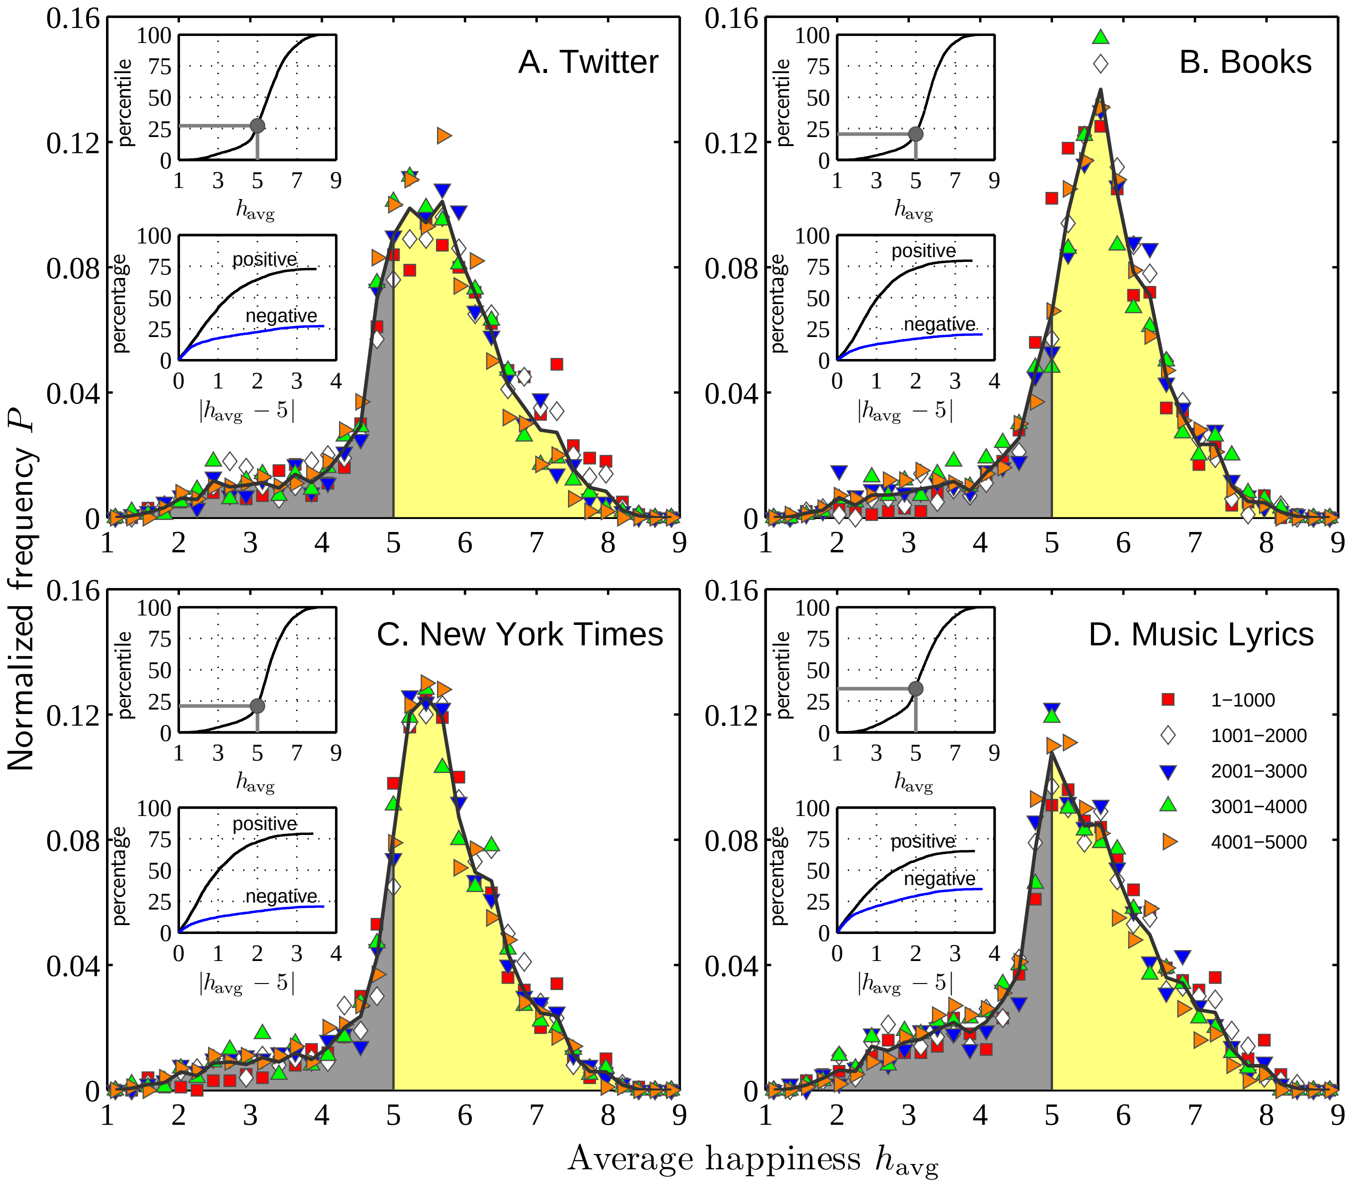

FIG. 1: Positivity bias in the English language: Normalized frequency distribution (solid black curves) of happiness scores for the 5000 most frequently used words in four corpora. Average happiness ratings for 10,222 words were obtained using Mechanical Turk with 50 evaluations per word for a total of 501,110 human evaluations (see main text). The yellow shade indicates words with average happiness scores above the neutral value of 5, gray those below. The symbols show normalized frequency distributions for subset word usage frequency ranges (see legend) demonstrating an internal scale-free consistency of positivity (see Fig. 2 for results of Kolmogorov-Smirnov tests). Upper inset plots show percentile locations and the lower inset plots show the number of words found when cumulating toward the positive and negative sides of the neutral score of 5. The distributions as shown were formed using 35 equal-sized bins; the number of bins does not change the visual form of the distributions appreciably, and an odd number ensures that the neutral score of 5 is a bin center. We employed binning only for visual display, using the raw data for all statistical analysis.

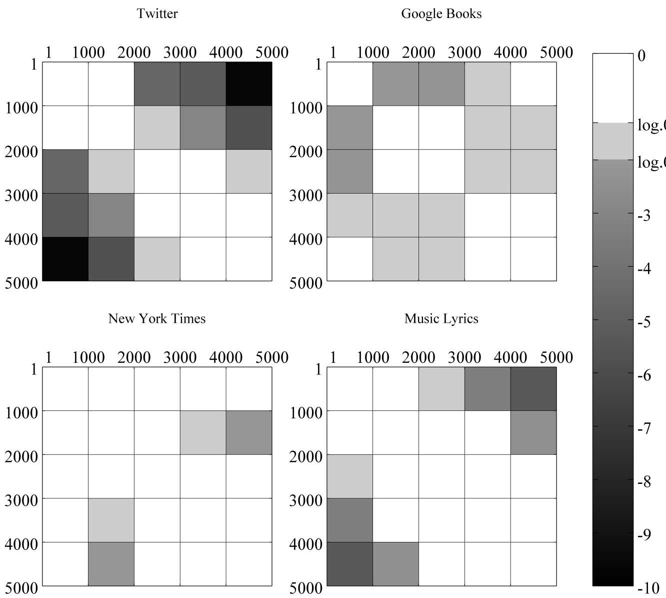

FIG. 2: Results of Kolmogorov-Smirnov tests comparing word happiness distributions shown in Fig. 1. For each corpus, the p-value reports the probability that the two samples being compared could come from the same distribution with lighter colors meaning more likely. The gray-scale corresponds to $\log_{10}(p\text{-value})$.

(GB), 78.38% (NYT), and 64.14% (ML). Equivalently, and as further supported by Fig. 1's upper inset plots of percentile location, we see the percentile corresponding to the neutral score of 5 is well below the median. The lower inset plots show how the number of positive and negative words increase as we cumulate moving away from the neutral score of 5; positive words are always more abundant further illustrating the positive bias. The mode average happiness of words is either above neutral (TW, GB, NYT) or located there (ML). Combining words across corpora, we see the same overall positivity bias for parts of speech, e.g., nouns and verbs (not shown), in agreement with previous work [1].

While the distributions do not match in detail across corpora, we do find they have an unexpected internal consistency with respect to usage frequency. In each plot in Fig. 1, we show distributions for subsets of 1000 words (symbols), ordered by frequency (e.g., 1–1000, 1001–2000, etc.), suggesting to the eye that common and rare words are similarly distributed in their perceived degree of positivity. Directly comparing usage frequency and average happiness, we measure Spearman's correlation coefficient to be all close to zero, (for TW, GB, NYT, and ML, r_s = 0.103, 0.013, 0.044, and 0.081), if positive and statistically significant (p-values = 2.4×10^{-13}, 3.5×10^{-1}, 2.0×10^{-3}, and 1.0×10^{-8}). In Fig. 2, we provide statistical evidence via p-values from Kolmogorov-Smirnov tests for each pairing of distributions. The three corpora NYT, ML, and GB show the most internal agreement, and we see in all corpora that neighboring ranges of 1000 frequencies could likely match in distribution. Of the 40 pair-wise comparisons, 29 show statistically significant matches ($p > 10^{-2}$).

Thus, while positive words outnumber negative words overall, there is no tendency for high frequency words to be more likely to be positive than low frequency words. The two aspects of positivity and usage frequency can only be separated with the kind of data we study here. Previous claims that positive words are used more frequently [1–3], suffered from insufficient, non-representative data. For example, Rozin et al. recently compared usage frequencies for just seven adjective pairs of positive-negative opposites [2]. Augustine et al. showed that average happiness and usage frequencies for 1034 words [14] were more positively correlated than we observe here [1]; however, since these words were chosen for their meaningful nature [14, 26, 27] rather than by their rate of occurrence.

In sum, our findings for these diverse English language corpora suggest that a positivity bias is universal, that the emotional spectrum of language is self-similar with respect to frequency, and that in our stories and writings we tend toward prosocial communication. Furthermore, a positivity bias is not inconsistent with many findings that negative emotions are more potent and diverse [28]. Our work calls for similar studies of other languages and dialects, examinations of corpora factoring in popularity (e.g., of books or articles), as well as investigations of other more specific emotional dimensions. Related work would explore changes in positivity bias over time, and correlations with quantifiable aspects of societal organization and function such as wealth, cultural norms, and political structures. Analyses of the emotional content of phrases and sentences in large-scale texts would also be a natural next, more complicated stage of research. Promisingly, we have shown elsewhere for Twitter that the average happiness of individual words correlates well with that of surrounding words in status updates [19].

[1] A. A. Augustine, M. R. Mehl, and R. J. Larsen, Social Psychological and Personality Science (2011).
[2] P. Rozin, L. Berman, and E. Royzman, Cognition and Emotion **24**, 536 (2010).
[3] J. Boucher and C. E. Osgood, Journal of Verbal Learning and Verbal Behavior **8**, 1 (1969).
[4] R. W. Schrauf and J. Sanchez, Journal of Multilingual and Multicultural Development **25**, 266 (2004).
[5] R. Axelrod, *The Evolution of Cooperation* (Basic Books, New York, 1984).
[6] M. A. Nowak, Science **314**, 1560 (2006).
[7] D. Ariely, *Predictably Irrational* (Harper Perennial, 2010).
[8] P. J. Richerson and R. Boyd, *Not by Genes Alone* (Uni-

versity of Chicago Press, Chicago, IL, 2005).

[9] D. Kahneman, J. Knetsch, and R. Thaler, Journal of Political Economy **98**, 1325 (1990).

[10] R. S. Nickerson, Review of General Psychology **2**, 175 (1998).

[11] E. Fehr and S. Gächter, Nature **415**, 137 (2002).

[12] T. Hobbes, *Leviathan* (Oxford University Press, Oxford, UK, 2009).

[13] A. Rand, *The Virtue of Selfishness* (New American Library, New York, NY, 1964).

[14] M. Bradley and P. Lang, Technical report C-1, University of Florida, Gainesville, FL (1999).

[15] C. Osgood, G. Suci, and P. Tannenbaum, *The Measurement of Meaning* (University of Illinois, Urbana, IL, 1957).

[16] A. Chmiel, J. Sienkiewicz, M. Thelwall, G. Paltoglou, K. Buckley, A. Kappas, and J. A. Hoyst, PLoS ONE **6**, e22207 (2011).

[17] R. Reisenzein, in *The structuralist program in psychology: Foundations and applications*, edited by W. Hans (Hogrefe & Huber, Gottingen, 1992), pp. 141–189.

[18] H. Kwak, C. Lee, H. Park, and S. Moon, in *Proceedings of the 19th international conference on World wide web* (ACM, New York, NY, USA, 2010), WWW '10, pp. 591–600, ISBN 978-1-60558-799-8.

[19] P. S. Dodds, K. D. Harris, I. M. Kloumann, C. A. Bliss, and C. M. Danforth (2011), `http://arxiv.org/abs/1101.5120`.

[20] `http://apiwiki.twitter.com/`.

[21] J.-B. Michel, Y. K. Shen, A. P. Aiden, A. Veres, M. K. Gray, The Google Books Team, J. P. Pickett, D. Hoiberg, D. Clancy, P. Norvig, et al., Science Magazine **331**, 176 (2011).

[22] `http://ngrams.googlelabs.com/`.

[23] E. Sandhaus, *The New York Times Annotated Corpus*, Linguistic Data Consortium, Philadelphia (2008).

[24] P. S. Dodds and C. M. Danforth, Journal of Happiness Studies (2009), doi:10.1007/s10902-009-9150-9.

[25] `https://www.mturk.com/`.

[26] A. Mehrabian and J. A. Russell, *An approach to environmental psychology* (MIT Press, Cambirdge, MA, 1974).

[27] F. S. Bellezza, A. G. Greenwald, and M. R. Banaji, Behavior Research Methods, Instruments & Computers **18**, 299 (1986).

[28] R. F. Baumeister, E. Bratslavsky, C. Finkenauer, and K. D. Vohs, Review of General Psychology **5**, 323 (2001).

The authors are indebted to conversations with B. Tivnan, N. Johnson, and Andrew Reece. The authors are grateful for the computational resources provided by the Vermont Advanced Computing Center which is supported by NASA (NNX 08A096G). KDH was supported by VT-NASA EPSCoR. PSD was supported by NSF CAREER Award # 0846668.